



Software Structures for Virtual Environments

Capps, Darken, Zyda
Naval Postgraduate School



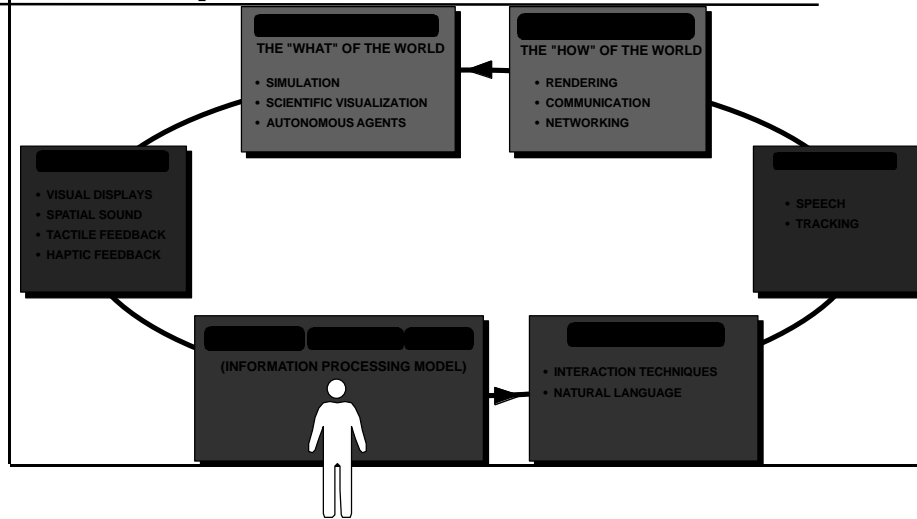
What does VR software look like?

One Thread, Multiple Threads

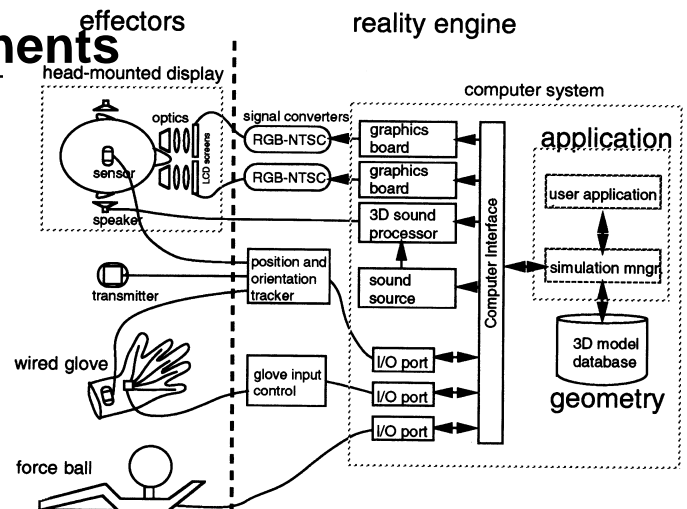
Important Subsystems

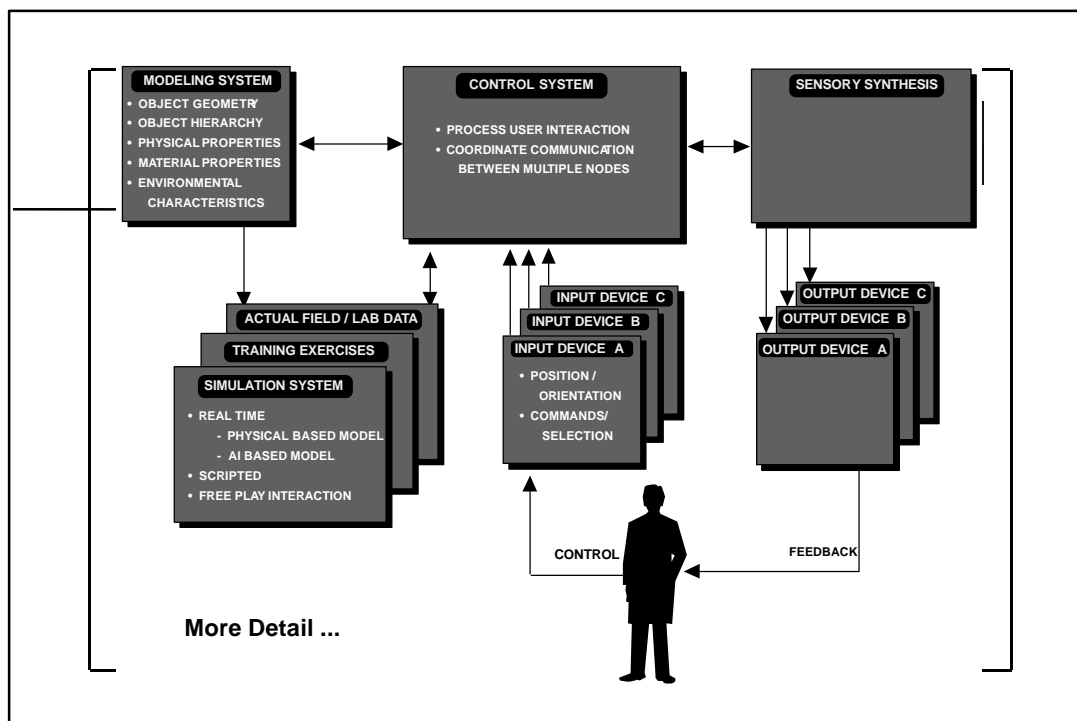
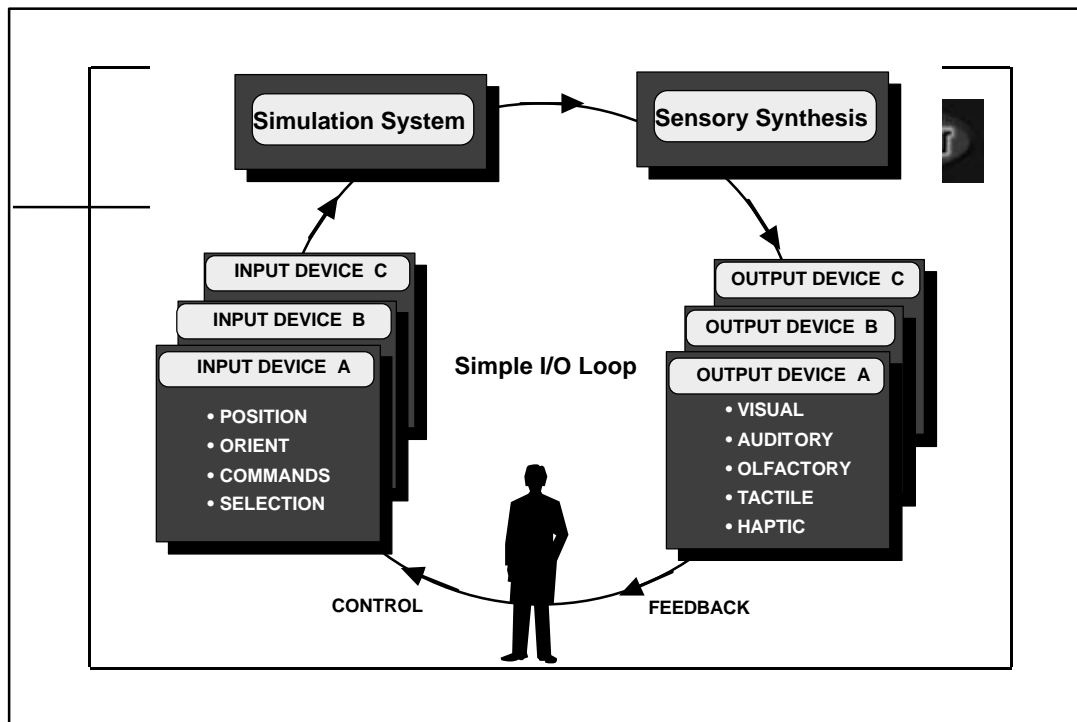
- Real-Time Rendering - Polygon Culling & Level of Detail Processing
- Real-Time Collision Detection and Response
- Computational Resource Management
- Interaction Management

Conceptual Model



System Components





Single Thread Networked VE

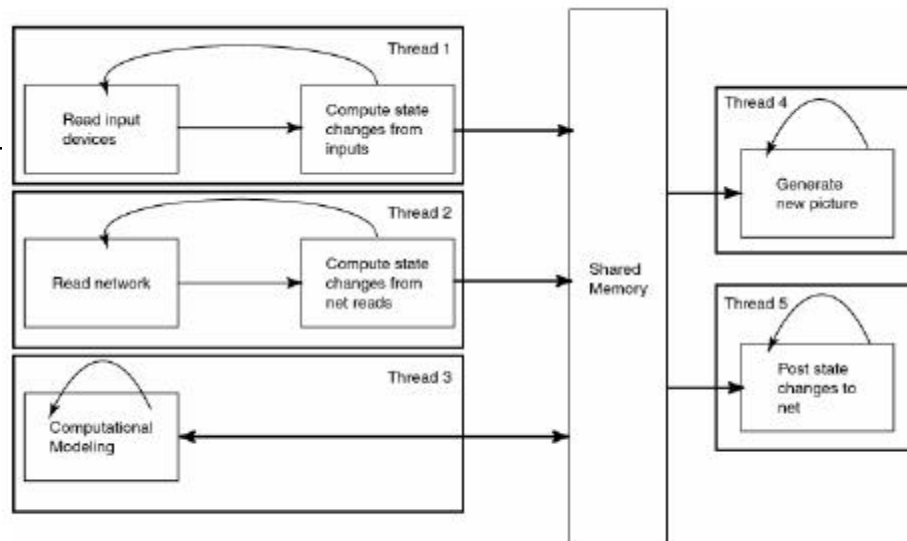
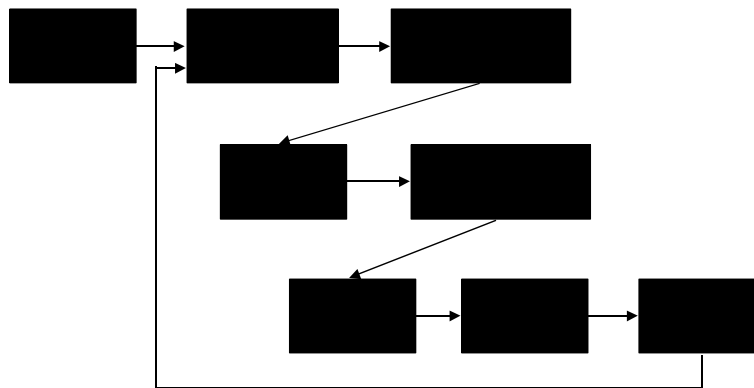


Figure 6-4 Multiple threads for the virtual world.

Real-Time Rendering - Polygon Culling & LODs



The “Generate New Picture” box on the previous slide is somewhat misleading.

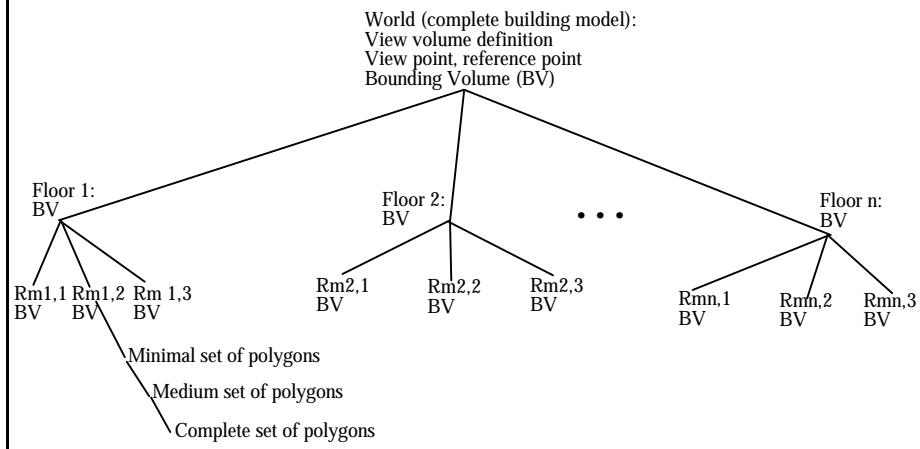
- It is not that simple unless we have unlimited graphics power. We usually don't.
- Most of the time we are trying to solve the Polygon Culling Problem.
 - Use available CPU cycles to cull most of our 3D model before we send it through the graphics pipeline ...

Hierarchical Data Structures for Polygon Culling



- The classic reference for this is the 1976 paper by James Clark “Hierarchical Geometric Models for Visible Surface Algorithms”.
- The idea in the Clark paper is to build a hierarchical data structure for the displayable world ...

Hierarchical Data Structure for the Displayable World



Part 1 of the Clark Paper



- We build a data structure that allows us to rapidly throw away most of the data for the world.
- We test the hierarchical bounding volumes to see if they are contained or partially contained in the current orientation of the view volume.
- If they are and are leaves in the tree, we draw them or continue our traversal.

Part 2 of the Clark Paper



- We only send the minimal required description of our object through the graphics pipeline.
- We use the distance from the viewer to the BV/object to determine which resolution of our objects to display, basically the pixel coverage of the final drawn object.
- This presupposes multiple resolution versions of each room.

View Volume

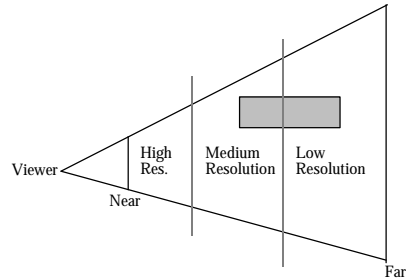


- We can decide rather quickly whether our BV is in the High, Medium or Low resolution sections.
 - Transform the BVs with the far clipping plane moved closer.
 - Clip into memory and make the decision (or perform this test in the CPU, if there are spare cycles).
 - We will end up with a list of BVs that should be displayed and a resolution at which they should be displayed.

View Volume



This the view volume cut into three regions or “levels of detail” (LOD).



How do we compute different LODs for our models?



This depends on the origins of our models...

- By hand - we do this by hand in our modeling tool, throwing away small polygons for the Low resolution versions of our models.
 - Some modeling tools will do this semi-automatically. They give you a cut at it and you can add polygons back in.

How do we compute different LODs for our models?

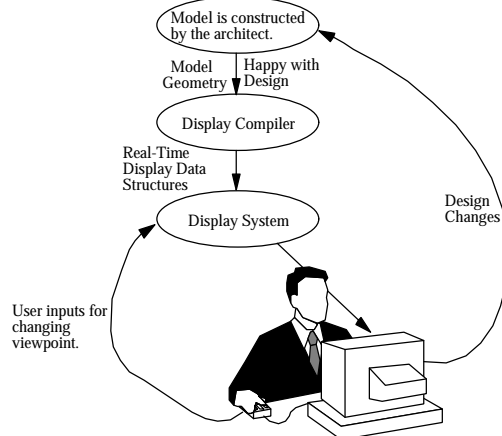


- Triangular decimation - a triangular mesh is fitted to your model and an appropriate algorithm is used to reduce the total number of triangles in the model.
 - There is some very nice work on this by Greg Turk and Hughe Hoppe and others...

Airey, Rohlf & Brooks Paper



- Precomputation of a hierarchical data structure for a building.



Airey, Rohlf & Brooks Paper Study of Architectural DBs



- The model is changed much less often than the viewpoint.
 - Means pre-processing the database (display compiler) is possible.
 - We could possibly make changes to the big data structure as we added new building components.
 - Perhaps the real-time data structure is updated in parallel. During the “think time” of the engineer at the workstation.

Airey, Rohlf & Brooks Paper Study of Architectural DBs



- Many buildings have high average depth complexity.
 - Any image computed from an interior viewpoint will have many surfaces covering each pixel.
 - Much of the model contributes NOTHING to any given image.

Airey, Rohlf & Brooks Paper Study of Architectural DBs



- Most polygons are axial.
 - They are parallel to 2 of the coordinate axes.
 - Most polygons are rectangles.

Airey, Rohlf & Brooks Paper Study of Architectural DBs



- The set of polygons that appears in each view changes slowly as the viewpoint moves.
 - Except when crossing certain thresholds.
Doors, windows --> portals.

Airey, Rohlf & Brooks Paper Study of Architectural DBs



- We can possibly have the notion of a “working set” of bounding volumes.
 - Based on a viewpoint.
 - Also see this in the Clark paper.
 - We could just incrementally add/subtract branches of our tree based on view point changes.

Airey, Rohlf & Brooks Paper Study of Architectural DBs

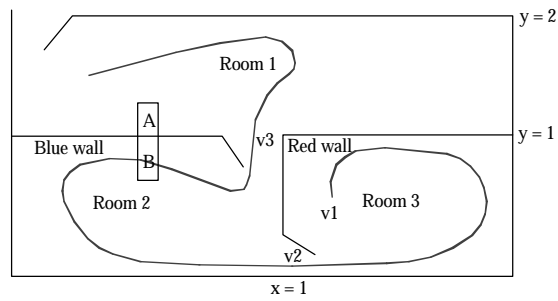


- This means that when we organize our data, our data inside one bounding volume should be co-located in CPU memory.
 - To make best use of the virtual memory system.

Portals and Viewpoints ...



- We have the notion of viewpoints at portals being indicators that we need to swap in major new blocks of data.



Airey, Rohlf & Brooks Paper Study of Architectural DBs



- Large planar surfaces are often structured into multiple, co-planar levels for modeling purposes, shading purposes and realism purposes.
 - Large walls that cross several rooms might be stored as multiple polygons of different color.
 - Perhaps BV-wise, we could use the larger wall better than the smaller components.
 - Notion of somehow taking advantage of such dividing planes in building our tree structure.

Airey, Rohlf & Brooks Paper Study of Architectural DBs



- For viewpoints inside the building, the role of surface interreflections in shading calculations is very important for spatial comprehension.
 - In the Airey system, there is an adaptive radiosity display algorithm.
 - When the viewpoint is not changing, the better radiosity view is displayed.

An adaptive system. Put up more detail when the system is not moving.

Model Space Subdivision



- UNC builds its data structures for display with a Display compiler.
 - Automatically subdivide their database into cells based on the union of “potentially visible sets” (PVS) for any viewpoint in a cell.
 - Viewpoint position tells which cells to display
 - That cell contains potentially visible information.

Model Space Subdivision



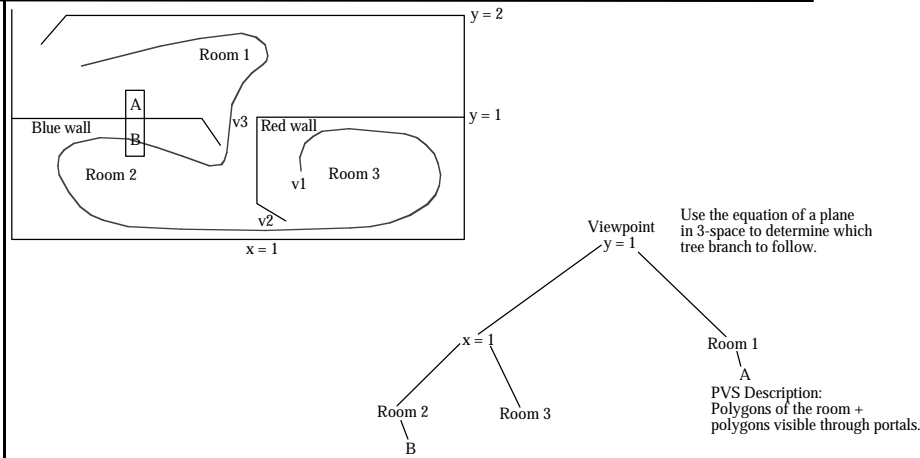
- A cell is a room plus any potentially visible polygons, polygons visible through portals.
- Computing PVSs is a hard problem.
- For any viewpoint, we must display the polygons for the room plus any possibly visible ones through any doors/portals.

Binary Space Partitioning



- Airey used a BSP-tree as the data structure for his building models.
- His paper describes how to automatically choose dividing planes.
- Use the biggest polygon, the ones with the most occlusion potential.

Binary Space Partitioning



Partition Priority for Polygons



- One of the key problems in model space subdivision is determining which are the best planes for splitting the geometric database.
- Airey came up with the idea of a “partition priority” for any polygon.

partition priority
for any polygon
(used to determine
best plane for splitting
the database.)

$$= 0.5 * \text{occlusion} + 0.3 * \text{balance} + 0.2 * \text{split}$$

The biggest known
polygons with best
occlusion potential
are weighted the
most.

BSP-tree
balance, i.e.
1/2 polygons
are on each
side of the
dividing
polygon.

Sometimes must
cut polygons by
the dividing plane
(want to minimize
that.)

Summary of Model Subdivision Results



\$	Polygons\$	Cells\$	Polys/ Cell\$	Polys/ Cell\$	Speed-up\$	Speed-up\$
\$	\$	\$	Average\$	Max.\$	Average\$	Min.\$
Sitt. Hall\$	7125\$	269\$	230\$	2195\$	30.98\$	3.25\$
Lobby\$	3949\$	54\$	466\$	2550\$	8.47\$	1.55\$
Church 1\$	7812\$	108\$	291\$	2055\$	26.85\$	3.8\$
Church 2\$	6037\$	16\$	1887\$	3477\$	3.2\$	1.74\$

Optimal Number of Polygons Per Cell?



- The number of polygons per cell is determined by:
 - the graphics hardware's fill capability
 - by the CPU capability to compute which cells to display

Papers Useful for Walkthrough



- (1) 1976 CACM Clark, James H. "Hierarchical Geometric Models for Visible Surface Algorithms"
- (2) Fuchs "Near-Real-Time Shaded Display of Rigid Objects" - BSP-tree fundamentals. SIGGRAPH ?
- (3) Brooks - 1986 Workshop on Interactive 3D Graphics - Early thoughts on walkthroughs.
"Walkthrough - A Dynamic Graphics System for Simulating Virtual Buildings"

Papers Useful for Walkthrough



- (4) Notes on the origin of BSP trees.
- (5) Airey, Rohlf & Brooks 1990 Symposium on Interactive 3D Graphics paper. "Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments"
- (6) Papers by Funkhouser
- (7) Paper by Teller & Sequin (SIGGRAPH 93)

Real-Time Collision Detection & Response



A Short Survey on Collision Detection

It's explored in the literature of:

- computer graphics
- robotics
- computational geometry
- computer animation
- physically-based modeling

Real-Time Collision Detection & Response



Numerous approaches:

- geometric reasoning[DK90]
- bounding volume hierarchy [Hub96]
- spatial representation [GASF94, NAT90]
- numerical methods[Cam97, GJK88]
- analytical methods[LM95, Sea93]

Real-Time Collision Detection & Response



However, many of these algorithms do not satisfy the demanding requirements of general-purpose collision detection for networked virtual environments.

Real-Time Collision Detection & Response



UNC has developed a mix of algorithms and systems for large interactive environments:

- I-COLLIDE [CLMP95]
- RAPID [GLM96]
- V-COLLIDE [HLC + 97].



How it Works

Brute force collision detection:

- $O(n^2)$ for convex polyhedron: compare each face on object A with each face on object B; then test if any point in A is inside of B or vice-versa.
- We do this process every frame!

A little better: spatial subdivision

- only compare two objects if they are in same region of space
- this isn't useful for a congested model!



Too slow?

Use axis-aligned bounding volumes:

- Sort in axial directions, and find pairs which overlap in all 3 dimensions

To save time, use spatial coherence: use a bounding volume big enough to hold the object at any orientation

- That way a spinning object does not require new calculation

Problem children



What about non-convex objects?

- Start with convex objects: divide or convex hull
- If convex versions collide, then must compare non-convex versions...
- ... then we're back to brute force

Real-Time Collision Detection & Response



***Public domain code is available for ftp at
<http://www.cs.unc.edu/~geom>***

***Netherlands research team has similar
system: **SOLID*****

<http://www.win.tue.nl/cs/tt/gino/solid/>.

Real-Time Collision Detection & Response



References

- [Cam97]** S. Cameron. *Enhancing gjk: Computing minimum and penetration distance between convex polyhedra. Proceedings of International Conference on Robotics and Automation, 1997.*
- [CLMP95]** J. Cohen, M. Lin, D. Manocha, and M. Ponamgi. *I-collide: An interactive and exact collision detection system for large-scale environments. In Proc. of ACM Interactive 3D Graphics Conference, pages 189{196, 1995.*

Real-Time Collision Detection & Response



- [DK90]** D. P. Dobkin and D. G. Kirkpatrick. *Determining the separation of pre-processed polyhedra { A uni ed approach. In Proc. 17th Internat. Colloq. Automata Lang. Program., volume 443 of Lecture Notes Comput. Sci., pages 400{413. Springer-Verlag, 1990.*
- [GASF94]** A. Garcia-Alonso, N. Serrano, and J. Flaquer. *Solving the collision detection problem. IEEE Comput. Graph. Appl., 14:36{43, May 1994.*

Real-Time Collision Detection & Response



[GJK88] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi. *A fast procedure for computing the distance between objects in three-dimensional space. IEEE J. Robotics and Automation*, vol RA-4:193{203, 1988.

[GLM96] S. Gottschalk, M. Lin, and D. Manocha. *Obb-tree: A hierarchical structure for rapid interference detection. In Proc. of ACM Siggraph'96*, pages 171{180, 1996.

Real-Time Collision Detection & Response



[HLC+97] T. Hudson, M. Lin, J. Cohen, S. Gottschalk, and D. Manocha. *V-collide: Accelerated collision detection for vrml. In Proc. of VRML Conference*, pages 119{125, 1997.

[Hub96] P. M. Hubbard. *Approximating polyhedra with spheres for time-critical collision detection. ACM Trans. Graph.*, 15(3):179{210, July 1996.